

IN THE SPECIFICATION

After the title, please add the following heading and paragraph:

“CROSS-REFERENCE TO RELATED APPLICATIONS

This application is a continuation of pending Application No. 09/684,330, filed October 10, 2000. Each and every document, including patents and publications, cited herein is incorporated herein by reference in its entirety as though recited in full.”

Please replace the paragraphs beginning on page 6, line 1 and ending on page 6, line 11 with the following:

“According to some embodiments, a document retrieval system receives information regarding the relevancy of documents retrieved in response to a current query. The retrieved documents are ranked in accordance with the relevancy information. A candidate query is formed from the rankings and an analysis of locations of the retrieved documents in a latent semantic index space formed from the retrieved documents. The candidate query is applied the document collection and documents retrieved in response to the candidate query are ranked in accordance with the received relevancy information. The ranking of documents retrieved in response to each query is compared, the query that produces the best ranking is chosen as the new query.

According to other embodiments for refining a strategy for ranking the results of a document retrieval operation, the results of a document retrieval process are indexed into a latent semantic index vector space. Information is received regarding relevancy of the results to an information need. The results are ranked based on the current ranking strategy. The current ranking strategy is modified based on the location of the ranked documents in the latent semantic

index vector space and the relevancy information. The results are ranked based on the new ranking strategy. The ranking strategy that produces results that better correspond to the relevancy information is chosen as the refined ranking strategy.

According to further embodiments, the invention includes methods for refining a current query. In these methods, a subset of documents retrieved in response to a current query is ranked with respect to relevancy of the retrieved documents to an information need. A subset of the documents retrieved in response to a candidate query is ranked with respect to the same standard. The ranking is a function of the context of the relevancy information and the location of each document in a latent semantic index vector space. The refined query is chosen as the query which produces a ranking indicating that the retrieved documents corresponding to the query more closely matches the relevancy information.

In yet other embodiments, the invention includes methods for forming a hypothesis for refining a query. In these embodiments, a document collection is indexed into a latent semantic vector space. Information regarding the relevancy with respect to information needs is received regarding the retrieved documents. At least one hypothesis is formed regarding modifications to a query based on analysis of the context of the relevancy information and the documents corresponding to the relevancy information in the vector space.”